COMP323 – Introduction to Computational Game Theory

Extensive Games with Perfect Information

Paul G. Spirakis

Department of Computer Science
University of Liverpool

Outline

Introduction

- The model of a strategic game suppresses the sequential structure of decision-making:
  - each decision-maker chooses her plan of action once and for all;
  - she is committed to this plan, which she cannot modify.
- The model of an extensive game allows us to study situations in which each decision-maker is free to change her mind as events unfold:
  - the sequential structure of decision-making is explicitly described.
- We will study a model in which each decision-maker is always fully informed about all previous actions.

The model

To describe an extensive game with perfect information, we need to specify:

- The set of players and their preferences (as for a strategic game);
- The order of the players' moves and the actions each player may take at each point, by specifying
  - the set of all sequences of actions that can possibly occur,
  - the player who moves at each point in each sequence.
- Each possible sequence of actions is a terminal history.
- The function that gives the player who moves at each point in each terminal history is the player function

Extensive game with perfect information

An extensive game has four components:

1. players;
2. terminal histories;
3. player function;
4. preferences for the players.

Before giving precise definitions, we will give an example illustrating them.

An example

Example

- An incumbent faces the possibility of entry by a challenger.
- The challenger may enter or not.
- If it enters, the incumbent may either acquiesce or fight.

We may model this situation as an extensive game with perfect information in which

- the terminal histories are (In, Acquiesce), (In, Fight), and Out;
- the player function assigns the challenger to the start of the game and the incumbent to the history In.

Actions

- At the start of an extensive game, and after any sequence of events, a player chooses an action.
- The sets of actions available to the players are not given explicitly.
- What is specified is the set of terminal histories and the player function, from which we can deduce the available sets of actions.

In the entry game example:

- the actions available to the challenger at the start of the game are In and Out, because these actions (and no others) begin terminal histories;
- the actions available to the incumbent are Acquiesce and Fight, because these actions (and no others) follow In in terminal histories.

Terminal histories

- The terminal histories of a game are specified as a set of sequences.
- Not every set of sequences is a legitimate set of terminal histories.
- If (C, D) is a terminal history, C should not be specified as a terminal history: after C is chosen at the start of the game, some player may choose D, so that the action C does not end the game.
- A sequence that is a proper subhistory of a terminal history cannot itself be a terminal history.
- This the only restriction we need to impose on a set of sequences so that the set be interpretable as a set of terminal histories.

Subhistories

Define the subhistories of a finite sequence $(a^1, a^2, \ldots, a^k)$ of actions to be

1. the empty history consisting of no actions, denoted $\emptyset$, and
2. all sequences of the form $(a^1, a^2, \ldots, a^m)$, where $1 \leq m \leq k$.

Similarly, define the subhistories of an infinite sequence $(a^1, a^2, \ldots)$ of actions to be

1. the empty history consisting of no actions, denoted $\emptyset$,
2. all sequences of the form $(a^1, a^2, \ldots, a^m)$, where $m \geq 1$, and
3. the entire sequence $(a^1, a^2, \ldots)$.

Subhistories

- A subhistory not equal to the entire sequence is called a proper subhistory.
- A sequence of actions that is a subhistory of some terminal history is called simply a history.

In the entry game example:

- The subhistories of (In, Acquiesce) are the empty history $\emptyset$ and the sequences In and (In, Acquiesce).
- The proper subhistories are the empty history and the sequence In.

Definition

Definition

An extensive game with perfect information consists of

- a set of players;
- a set of sequences (terminal histories) with the property that no sequence is a proper subhistory of any other sequence;
- a function (the player function) that assigns a player to every sequence that is a proper subhistory of some terminal history;
- for each player, preferences over the set of terminal histories.

- The set of terminal histories is the set of all sequences of actions that may occur.
- The player assigned by the player function to any history h is the player who takes an action after h.
- We may specify a player's preferences by giving a payoff function that represents them.

Example: entry game

Suppose that

- the best outcome for the challenger is that it enters and the incumbent acquiesces;
- the worst outcome for the challenger is that it enters and the incumbent fights;
- the best outcome for the incumbent is that the challenger stays out; and
- the worst outcome for the incumbent is that it enters and there is a fight.

Then the situation may be modeled as an extensive game with perfect information.

Example: entry game

- **Players** The challenger and the incumbent.
- **Terminal histories** $(In, Acquiesce)$, $(In, Fight)$, and Out.
- **Player function:** $P(\emptyset) = Challenger$ and $P(In) = Incumbent$.
- **Preferences:**
  - Challenger's preferences are represented by the payoff function u_1:

$$u_1(In, Acquiesce) = 2 \quad u_1(Out) = 1 \quad u_1(In, Fight) = 0 \ .$$

  - Incumbent's preferences are represented by the payoff function u_2:

$$u_2(In, Acquiesce) = 1 \quad u_2(Out) = 2 \quad u_2(In, Fight) = 0 \ .$$

Representation of an extensive game

The entry game is readily illustrated in a diagram:

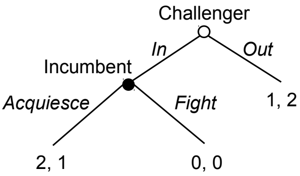

- The small circle at the top represents the empty history (the start of the game).
- The label above a node indicatesthe player who chooses an action.
- The branches represent the player's choices.
- The pair of numbers beneath each terminal history gives the players' payoffs to that history.

Actions

- The sets of actions available to the players at their various moves are not directly specified.
- These can be deduced from the set of terminal histories and the player function.
- If, for some nonterminal history h, the sequence (h, a) is a history, then a is one of the actions available to the player who moves after h.
- Thus the set of all actions available to the player who moves after h is

$$A(h) = \{a : (h, a) \text{ is a history}\} \ .$$

Actions

For example, for the entry game:

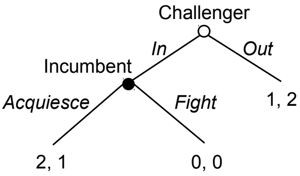

- The histories are $\emptyset$, In, Out, $(In, Acquiesce)$, and $(In, Fight)$.
- The set of actions available to the challenger who moves at the start of the game is $A(\emptyset) = \{In, Out\}$.
- The set of actions available to the incumbent who moves after the history In is $A(In) = \{Acquiesce, Fight\}$.

Finiteness

- Terminal histories are allowed to be infinitely long.
- If the length of the longest terminal history is in fact finite, we say that the game has a finite horizon.
- A game with a finite horizon may have infinitely many terminal histories: some player might have infinitely many actions after some history.
- If a game has a finite horizon and finitely many terminal histories we say it is finite.
- A game that is not finite cannot be represented in a diagram!

Perfect information

An extensive game with perfect information models a situation in which each player, when choosing an action,

1. knows all actions chosen previously (has perfect information), and
2. always moves alone (rather than simultaneously with other players).

The model encompasses several situations:

- A race (e.g., between firms developing a new technology) is modeled as an extensive game in which the parties alternately decide how much effort to expend.
- Parlor games such as chess, in which there are no random events, the players move sequentially, and each player always knows all actions taken previously, may also be modeled as extensive games with perfect information.

Backward induction

In the entry game:

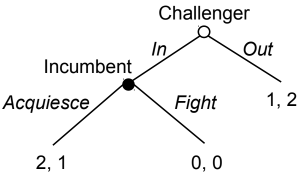

- The challenger will enter and the incumbent will subsequently acquiesce.
- The challenger can reason that if it enters then the incumbent will acquiesce, because doing so is better for the incumbent than fighting.
- Given that the incumbent will respond to entry in this way, the challenger is better off entering.

Backward induction

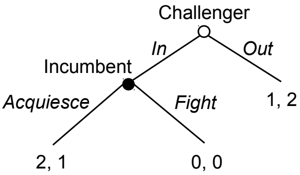

This line of argument is called backward induction:

A player who has to move deduces, for each of her possible actions, the actions that the players (including herself) will subsequently rationally take, and chooses the action that yields the terminal history she most prefers.

Backward induction

Backward induction cannot be applied to every extensive game with perfect information:

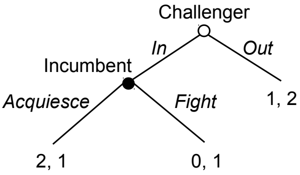

- If the challenger enters, the incumbent is indifferent between acquiescing and fighting.
- Backward induction does not tell the challenger what the incumbent will do in this case.
- Games with infinitely long histories present another difficulty for backward induction.

Nash equilibrium

Another approach to defining equilibrium takes off from the notion of Nash equilibrium:

- It seeks to model patterns of behavior that can persist in a steady state.
- The resulting notion of equilibrium applies to all extensive games with perfect information.
- In games in which backward induction is well-defined, this approach turns out to lead to the backward induction outcome, so that there is no conflict between the two ideas.

Strategies

A player's strategy specifies the action the player chooses for every history after which it is her turn to move.

Definition

A strategy of player i in an extensive game with perfect information is a function that assigns to each history h after which it is player i's turn to move (i.e. $P(h) = i$, where P is the player function) an action in $A(h)$ (the set of actions available after h).

Strategies: an example

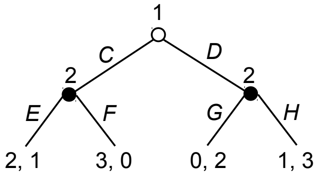

- Player 1 moves only at the start of the game (i.e. after the empty history), when the actions available to her are C and D.
- Thus she has two strategies: one that assigns C to the empty history, and one that assigns D to the empty history.

Strategies: an example

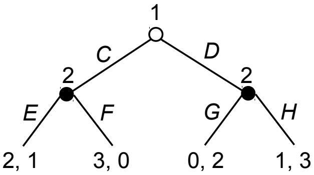

- Player 2 moves after both the history C and the history D.
- After C the actions available to her are E and F, and after D the actions available to her are G and H
- Thus a strategy of player 2 is a function that assigns either E or F to the history C, and either G or H to the history D.
- That is, player 2 has four strategies: EG, EH, FG, FH.

Strategies: an example

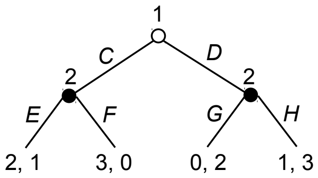

- Each of player 2's strategies may be interpreted as a plan of action or contingency plan: it specifies what player 2 does if player 1 chooses C, and what she does if player 1 chooses D.
- A player's strategy provides sufficient information to determine her plan of action: the actions she intends to take, whatever the other players do.

Outcomes

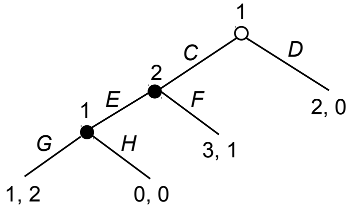

- The outcome of the strategy pair (DG, E) is the terminal history D.
- The outcome of (CH, E) is the terminal history (C, E, H).
- Note that the outcome $O(s)$ of the strategy profile s depends only on the players' plans of action, not their full strategies.
- To determine $O(s)$ we do not need to refer to any component of any player's strategy that specifies her actions after histories precluded by that strategy.

Nash equilibrium

- Nash equilibrium: a strategy profile from which no player wishes to deviate, given the other players' strategies.
- One way to find the Nash equilibria of an extensive game in which each player has finitely many strategies is to
  1. list each player's strategies,
  2. find the outcome of each strategy profile, and
  3. analyze this information as for a strategic game.

Strategic form

We construct the following strategic game, known as the strategic form of the extensive game:

- Players: the set of players in the extensive game.
- Actions: Each player's set of actions is her set of strategies in the extensive game.
- Preferences: Each player's payoff to each action profile is her payoff to the terminal history generated by that action profile in the extensive game.

The set of Nash equilibria of any extensive game with perfect information is the set of Nash equilibria of its strategic form.

Example: the entry game

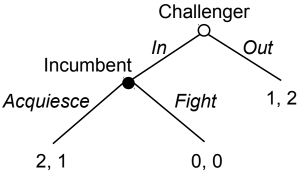

- the challenger has two strategies, *In* and *Out*
- the incumbent has two strategies, *Acquiesce* and *Fight*.

Strategic form of the game:

		Incumbent	
		Acquiesce	*Fight*
Challenger	*In*	(2,1)	(0,0)
	Out	(1,2)	(1,2)

Example: the entry game

		Incumbent	
		Acquiesce	*Fight*
Challenger	*In*	(2,1)	(0,0)
	Out	(1,2)	(1,2)

Two Nash equilibria: $(In, Acquiesce)$ and $(Out, Fight)$

- Equilibrium $(In, Acquiesce)$ is the pattern of behavior isolated by backward induction.
- In equilibrium $(Out, Fight)$
  - the challenger always chooses *Out*; this strategy is optimal given the incumbent's strategy to fight in the event of entry;
  - incumbent's strategy *Fight* is optimal given the challenger's strategy,

  thus neither player can increase its payoff by choosing a different strategy, *given the other player's strategy*.

Subgame perfect equilibrium

- The notion of Nash equilibrium ignores the sequential structure of an extensive game.
- It treats strategies as choices made once and for all before play begins.

Subgame perfect equilibrium is a notion of equilibrium that models a robust steady state.

- Each player's strategy is optimal, given the other players' strategies, not only at the start of the game, but after every possible history.

Subgames

We first define the notion of a subgame:

- For any nonterminal history h, the subgame following h is the part of the game that remains after h has occurred.
- For example, the subgame following the history *In* in the entry game is the game in which the incumbent is the only player, and there are two terminal histories, *Acquiesce* and *Fight*.

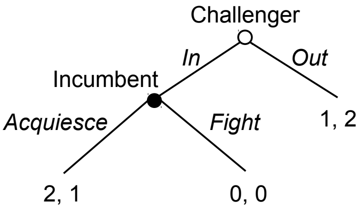

Subgames

- The subgame following the empty history $\emptyset$ is the entire game.
- Every other subgame is called a proper subgame.
- Because there is a subgame for every nonterminal history, the number of subgames is equal to the number of nonterminal histories.

Example:

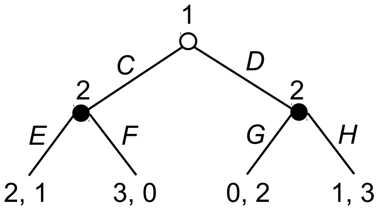

The above game has three nonterminal histories (the empty history, C, and D), and hence three subgames: the whole game (the part of the game following the empty history), the game following the history C, and the game following the history D.

Subgame perfect equilibrium

- In an equilibrium that corresponds to a perturbed steady state in which every history sometimes occurs, the players' behavior must correspond to a steady state in every subgame, not only in the whole game.

Definition

A subgame perfect equilibrium is a strategy profile s^* with the property that in no subgame can any player i do better by choosing a strategy different from s_i^*, given that every other player j adheres to s_j^*,

Subgame perfect equilibrium

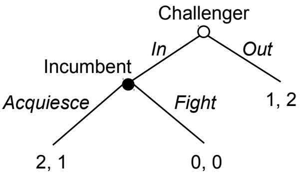

The Nash equilibrium $(Out, Fight)$ of the entry game is not a subgame perfect equilibrium:

- in the subgame following the history *In*, the strategy *Fight* is not optimal for the incumbent, since the incumbent is better off choosing *Acquiesce* than it is choosing *Fight*

The Nash equilibrium $(In, Acquiesce)$ is a subgame perfect equilibrium:

- each player's strategy is optimal, given the other players strategy, both in the whole game, and in the subgame following the history *In*.

Subgame perfect equilibrium and Nash equilibrium

- In a subgame perfect equilibrium every player's strategy is optimal, in particular, after the empty history. Thus:

Every subgame perfect equilibrium is a Nash equilibrium.

- A subgame perfect equilibrium generates a Nash equilibrium in every subgame.
- Further, any strategy profile that generates a Nash equilibrium in every subgame is a subgame perfect equilibrium.

Definition

A subgame perfect equilibrium is a strategy profile that induces a Nash equilibrium in every subgame.

Example: the entry game

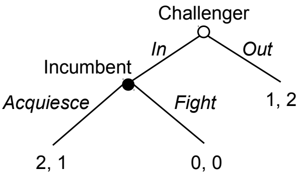

The Nash equilibrium (*In*, *Acquiesce*) is a subgame perfect equilibrium because

1. it is a Nash equilibrium, so that at the start of the game the challenger's strategy *In* is optimal, given the incumbent's strategy *Acquiesce*, and
2. after the history *In*, the incumbent's strategy *Acquiesce* in the subgame is optimal.

Example: variant of the entry game

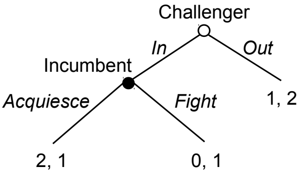

Two Nash equilibria, (*In*, *Acquiesce*) and (*Out*, *Fight*).

- Both of these equilibria are subgame perfect equilibria, because after the history *In* both *Fight* and *Acquiesce* are optimal for the incumbent.

Interpretation of subgame perfect equilibria

- A subgame perfect equilibrium of an extensive game corresponds to a steady state in which all players, on rare occasions, take nonequilibrium actions, so that after long experience each player forms correct beliefs about the other players' entire strategies, and thus knows how the other players will behave in every subgame.
- Given these beliefs, no player wishes to deviate from her strategy either at the start of the game or after any history.
- This interpretation does not require a player to know the other players' preferences, or to think about the other players' rationality.
- It entails interpreting a strategy as a plan specifying a player's actions not only after histories consistent with the strategy, but also after histories that result when the player chooses arbitrary alternative actions, perhaps because she makes mistakes or deliberately experiments.

Finite horizon games

- We can find the subgame perfect equilibria by finding the Nash equilibria and checking whether each of these equilibria is subgame perfect.
- In a game with a finite horizon the set of subgame perfect equilibria may be found more directly by using an extension of the procedure of backward induction.

Backward induction

The length of a subgame is the length of the longest history in the subgame.

The procedure of backward induction works as follows:

1. We start by finding the optimal actions of the players who move in the subgames of length 1 (the “last” subgames).
2. Taking these actions as given, we find the optimal actions of the players who move first in the subgames of length 2.
3. We continue working back to the beginning of the game, at each stage k finding the optimal actions of the players who move at the start of the subgames of length k, given the optimal actions we have found in all shorter subgames.

Backward induction: an example

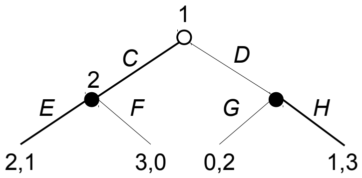

Subgames of length 1: The game has two such subgames, in both of which player 2 moves.

- In the subgame following C, player 2's optimal action is E.
- In the subgame following D, her optimal action is H.

Backward induction: an example

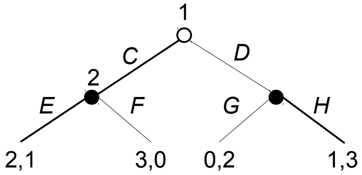

Subgames of length 2: The game has one such subgame, namely the entire game, at the start of which player 1 moves. Given the optimal actions in the subgames of length 1:

- player 1's choosing C at the start of the game yields her a payoff of 2, whereas
- her choosing D yields her a payoff of 1.

Thus player 1's optimal action at the start of the game is C.

$\Longrightarrow$ The game has no subgame of length greater than 2, so the procedure of backward induction yields the strategy pair (C, EH).

Extension of backward induction

- In any game in which the procedure selects a single action for the player who moves at the start of each subgame, the strategy profile selected is the unique subgame perfect equilibrium of the game (a complete proof is not trivial!).
- What happens in a game in which at the start of some subgames more than one action is optimal?
- An extension of the procedure of backward induction locates all subgame perfect equilibria.
- This extension traces back separately the implications for behavior in the longer subgames of every combination of optimal actions in the shorter subgames.

Procedure of backward induction

Backward induction may be described compactly for an arbitrary game as follows.

Step 1.

- Find, for each subgame of length 1, the set of optimal actions of the player who moves first.
- Index the subgames by j, and denote by $S_j^*(1)$ the set of optimal actions in subgame j.
  (If the player who moves first in subgame j has a unique optimal action, then $S_j(1)$ contains a single action.)

Procedure of backward induction

Step 2.

- For each combination of actions consisting of one from each set $S_j^*(1)$, find, for each subgame of length two, the set of optimal actions of the player who moves first.
- The result is a set of strategy profiles for each subgame of length two.
- Denote by $S_\ell^*(2)$ the set of strategy profiles in subgame ℓ.

Procedure of backward induction

Step 3, 4, ...

- Continue by examining successively longer subgames until reaching the start of the game.
- At each stage k, for each combination of strategy profiles consisting of one from each set $S_p^*(k-1)$ constructed in the previous stage, find, for each subgame of length k, the set of optimal actions of the player who moves first, and hence a set of strategy profiles for each subgame of length k.

Subgame perfect equilibrium and backward induction

The set of strategy profiles that the procedure of backward induction yields for the whole game is the set of subgame perfect equilibria of the game.

Proposition

The set of subgame perfect equilibria of a finite horizon extensive game with perfect information is equal to the set of strategy profiles isolated by the procedure of backward induction.

Note: A complete proof is not trivial.

Existence of subgame perfect equilibrium

Proposition

Every finite extensive game with perfect information has a subgame perfect equilibrium.

Proof.

- A finite game not only has a finite horizon, but also a finite number of terminal histories.
- The player who moves first in any subgame has finitely many actions; at least one action is optimal.
- Thus in such a game the procedure of backward induction isolates at least one strategy profile.
- Using the Proposition stated before, we conclude that every finite game has a subgame perfect equilibrium. □

Existence of subgame perfect equilibrium

- This existence result does not claim that a finite extensive game has a single subgame perfect equilibrium.
- A finite horizon game in which some player does not have finitely many ac- tions after some history may or may not possess a subgame perfect equilibrium.

A simple example of a game that does not have a subgame perfect equilibrium:

- Consider the trivial game in which a single player chooses a number *less than* 1 and receives a payoff equal to the number she chooses.
- There is no greatest number less than one, so the single player has no optimal action.
- Thus the game has no subgame perfect equilibrium.

The ultimatum game

Bargaining over the division of a pie may naturally be modeled as an extensive game:

Two people use the following procedure to split \$$c$:

- Person 1 offers person 2 an amount of money up to \$$c$.
- If 2 accepts this offer then 1 receives the remainder of the \$$c$.
- If 2 rejects the offer then neither person receives any payoff.
- Each person cares only about the amount of money she receives, and (naturally!) prefers to receive as much as possible.
- Assume that the amount person 1 offers can be any number, not necessarily an integral number of cents.
- Then the procedure can be modeled by an extensive game, known as the ultimatum game.

The ultimatum game

- Players: The two people.
- Terminal histories: The set of sequences (x, Z), where x is a number with $0 \leq x \leq c$ and Z is either Y ("yes, I accept") or N ("no, I reject").
- Player function: $P(\emptyset = 1$ and $P(x) = 2$ for all x.
- Preferences: Preferences are represented by payoffs equal to the amounts of money one receives.
  - For the terminal history (x, Y) person 1 receives $c - x$ and person 2 receives x.
  - For the terminal history (x, N) each person receives 0.

The ultimatum game

The game has a finite horizon, so we can use backward induction to find its subgame perfect equilibria.

Subgames of length 1: person 2 either accepts or rejects an offer of person 1.

- For every possible offer of person 1, there is such a subgame.
- In the subgame that follows an offer x of person 1 for which $x > 0$, person 2's optimal action is to accept (if she rejects, she gets nothing).
- In the subgame that follows the offer $x = 0$, person 2 is indifferent between accepting and rejecting.
- Thus in a subgame perfect equilibrium person 2's strategy either accepts all offers (including 0), or accepts all offers $x > 0$ and rejects the offer $x = 0$.

The ultimatum game

Now consider the **whole game**: for each possible subgame perfect equilibrium strategy of person 2, we need to find the optimal strategy of person 1.

- If person 2 accepts all offers (including 0), then person 1's optimal offer is 0 (which yields her the payoff c).
- If person 2 accepts all offers except zero, then no offer of person 1 is optimal!
  - No offer $x > 0$ is optimal, because the offer x/2 (for example) is better, given that person 2 accept both offers.
  - An offer of 0 is not optimal because person 2 rejects it, leading to a payoff of 0 for person 1, who is thus better off offering any positive amount less than c.

Conclusion: The only subgame perfect equilibrium of the game is the strategy pair in which person 1 offers 0 and person 2 accepts all offers. In this equilibrium, person 1's payoff is c and person 2's payoff is zero.

The holdup game

- Before engaging in an ultimatum game in which she may accept or reject an offer of person 1, person 2 takes an action that affects the size c of the pie to be divided.
- She may exert little effort, resulting in a small pie, of size c_L, or great effort, resulting in a large pie, of size c_H
- She dislikes exerting effort.
- Specifically, assume that her payoff is xE if her share of the pie is x, where $E = L$ if she exerts little effort and $E = H > L$ if she exerts great effort.
- The extensive game that models this situation is known as the holdup game.

The holdup game
Subgame perfect equilibrium

- Each subgame that follows person 2's choice of effort is an ultimatum game, and thus has a unique subgame perfect equilibrium, in which person 1 offers 0 and person 2 accepts all offers.
- Now consider person 2's choice of effort at the start of the game.
  - If she chooses L then her payoff, given the outcome in the following subgame, is L.
  - If she chooses H then her payoff is H.

  Consequently she chooses L
- Thus the game has a unique subgame perfect equilibrium, in which person 2 exerts little effort and person 1 obtains all of the resulting small pie.

The holdup game

Subgame perfect equilibrium

- This equilibrium does not depend on the values of c_L, c_H, L, and H (given that $H > L$).
- Even if c_H is much larger than c_L, but H is only slightly larger than L, person 2 exerts little effort in the equilibrium, although both players could be much better off if person 2 were to exert great effort and person 2 were to obtain some of the extra pie.
- No such superior outcome is sustainable in an equilibrium because person 2, having exerted great effort, may be "held up" for the entire pie by person 1.

Stackelberg's duopoly game

Consider a market in which there are two firms, both producing the same good.

- Firm i's cost of producing q_i units of the good is $C_i(q_i)$.
- The price at which output is sold when the total output is Q is $P_d(Q)$.
- Each firm's strategic variable is output, but the firms make their decisions sequentially, rather than simultaneously: one firm chooses its output, then the other firm does so, knowing the output chosen by the first firm.

This situation can be modeled by the Stackelberg's duopoly game.

Stackelberg's duopoly game

General model

- Players: The two firms.
- Terminal histories: The set of all sequences (q_1, q_2) of outputs for the firms (each q_i is a nonnegative number).
- Player function: $P(\emptyset) = 1$ and $P(q_1) = 2$ for all q_1.
- Preferences: The payoff of firm i to the terminal history (q_1, q_2) is its profit

$$q_i P_d(q_1 + q_2) - C_i(q_i) \quad \text{for } i = 1, 2 \ .$$

Note:

- Firm 1 moves at the start of the game, thus a strategy of firm 1 is simply an output.
- Firm 2 moves after every history in which firm 1 chooses an output, thus a strategy of firm 2 is a function that associates an output for firm 2 with each possible output of firm 1.

Stackelberg's duopoly game

Backwards induction

The game has a finite horizon, so we may use backward induction:

- Suppose that for each output q_1 of firm 1 there is one output $b_2(q_1)$ of firm 2 that maximize its profit. Then in any subgame perfect equilibrium, firm 2's strategy is b_2.
- *Given the strategy of firm 2*, when firm 1 chooses the output q_1, firm 2 chooses the output $b_2(q_1)$, resulting in a total output of $q_i + b_2(q_1)$, and a price of $P_d(q_1 + b_2(q_1))$. Thus firm 1's output in a subgame perfect equilibrium is a value of q_1 that maximizes

$$q_1 P_d(q_1 + b_2(q_1)) - C_1(q_1) \ .$$

Suppose that there is one such value of q_1, denote q_1^*.

Stackelberg's duopoly game

Backwards induction

We conclude:

- If firm 2 has a unique best response $b_2(q_1)$ to each output q_1, and firm 1 has a unique best action q_1^*, given firm 2's best responses, then the subgame perfect equilibrium of the game is $q_1^*, b_2)$.
- The output chosen by firm 2, given firm 1's equilibrium strategy, is $q_2^* = b_2(q_1^*)$.
- When firm 1 chooses any output q_1, the outcome, given that firm 2 uses its equilibrium strategy, is the pair of outputs $(q_1, b_2(q_1))$.
- As firm 1 varies its output, the outcome varies along firm 2's best response function b_2
- Thus we can characterize the subgame perfect equilibrium outcome (q_1^*, q_2^*) as the point on firm 2's best response function that maximizes firm 1's profit.

Stackelberg's duopoly game

Example: constant unit and linear inverse demand

Suppose that $C_i(q_i) = cq_i$ for $i = 1, 2$, and

$$P_d(Q) = \begin{cases} a - Q & \text{if } Q \leq a \\ 0 & \text{if } Q > a \end{cases},$$

where $0 < c < a$.

- Firm 2 has a unique best response to each output q_1 of firm 1:

$$b_2(q_1) = \begin{cases} \frac{1}{2}(a - c - q_1) & \text{if } q_1 \leq a - c \\ 0 & \text{if } q_1 > a - c \end{cases}.$$

- Thus in a subgame perfect equilibrium firm 2's strategy is b_2 and firm 1's strategy is the output q_1 that maximizes

$$q_1(a - c - (q_1 + 1/2(a - c - q_1))) = 1/2q_1(a - c - q_1) \ .$$

- Its maximizer is $q_1 = 1/2(a - c)$.

Stackelberg's duopoly game

Example: constant unit and linear inverse demand

We conclude:

- The game has a unique subgame perfect equilibrium, in which firm's 1 strategy is the output $1/2(a-c)$ and firm's 2 strategy is b_2.
- The outcome of the equilibrium is that firm 1 produces the output

$$q_1^* = 1/2(a-c)$$

  and frim 2 produces the output

$$q_2^* = b_2(q_1^*) = b_2(1/2(a-c)) = 1/4(a-c) \ .$$

- Firm 1's profit is

$$q_1^*(P_d(q_1^* + q_2^*) - c) = 1/8(a-c)^2$$

  and firm 2's profit is

$$q_2^*(P_d(q_1^* + q_2^*) - c) = 1/16(a-c)^2 \ .$$

Buying votes

- A legislature has k members, where k is an odd number.
- Two rival bills, X and Y, are being considered.
- The bill that attracts the votes of a majority of legislators will pass.
- Interest group X favors bill X, whereas interest group Y favors bill Y.
- Each group wishes to entice a majority of legislators to vote for its favorite bill.
- First interest group X gives an amount of money (possibly zero) to each legislator, then interest group Y does so. Each interest group wishes to spend as little as possible.
- Group X values the passing of bill X at $V_X > 0$ and the passing of bill Y at zero, and group Y values the passing of bill Y at $V_Y > 0$ and the passing of bill X at zero.

Buying votes

- Each legislator votes for the favored bill of the interest group that offers her the most money; a legislator to whom both groups offer the same amount of money votes for bill Y (an arbitrary simplifying assumption).
- For example, if $k = 3$, the amounts offered to the legislators by group X are $x = (100, 50, 0)$, and the amounts offered by group Y are $y = (100, 0, 50)$, then legislators 1 and 3 vote for Y and legislator 2 votes for X, so that Y passes.
- This situation can be modeled as an extensive game with perfect information.

Buying votes

Extensive game

- Players: The two interest groups, X and Y.
- Terminal histories: The set of all sequences (x, y), where x is a list of payments to legislators made by interest group X and y is a list of payments to legislators made by interest group Y (x and y are lists of k nonnegative integers).
- Player function: $P(\emptyset) = X$ and $P(x) = Y$ for all x.
- Preferences: The preferences of interest group X are represented by the payoff function

$$\begin{cases} V_X - (x_1 + \cdots + x_k) & \text{if bill } X \text{ passes} \\ -(x_1 + \cdots + x_k) & \text{if bill } Y \text{ passes,} \end{cases}$$

where bill Y passes after the terminal history (x, y) iff the number of components of y that are at least equal to the corresponding components of x is at least $1/2(k + 1)$ (a bare majority). The preferences of Y are represented by the analogous function.

Buying votes

Example 1

Suppose that $k = 3$ and $V_X = V_Y = 300$.

- The most group X is willing to pay to get bill X passed is 300.
- For any payments it makes to the three legislators that sum to at most 300, two of the payments sum to at most 200, so that if group Y matches these payments it spends less than $V_Y = 300$ and gets bill Y passed.
- Thus in any subgame perfect equilibrium group X makes no payments, group Y makes no payments, and (given the tie-breaking rule) bill Y is passed.

Buying votes

Example 2

Now suppose $k = 3$, $V_X = 300$, and $V_Y = 100$.

- By paying each legislator more than 50, group X makes matching payments by group Y unprofitable: only by spending more than $V_Y = 100$) can group Y cause bill Y to be passed.
- However, there is no subgame perfect equilibrium in which group X pays each legislator more than 50, because it can always pay a little less and still prevent group Y from profitably matching.
- In the only subgame perfect equilibrium group X pays each legislator exactly 50, and group Y makes no payments. Given group X's action, group Y is indifferent between matching X's payments (so that bill Y is passed), and making no payments.
- However, there is no subgame perfect equilibrium in which group Y matches group X's payments, because then group X could increase its payments a little, making matching payments by group Y unprofitable.

Buying votes

Subgame perfect equilibria

For arbitrary values of the parameters the subgame perfect equilibrium outcome takes one of the forms in these two examples:

1. either no payments are made and bill Y is passed, or
2. group X makes payments that group Y does not wish to match, group Y makes no payments, and bill X is passed.

To find the subgame perfect equilibria in general, we may use backward induction.

Buying votes

Subgame perfect equilibria

First consider group Y's best response to an arbitrary strategy x of group X.
Let $\mu = 1/2(k+1)$ and denote by m_x the sum of the smallest μ components of x.

- If $m_x < V_Y$ then group Y can buy off a bare majority of legislators for less than V_Y, so that its best response to x is to match group X's payments to the μ legislators to whom group X's payments are smallest. The outcome is that bill Y is passed.
- If $m_x > V_Y$ then the cost to group Y of buying off any majority of legislators exceeds V_Y, so that group Y's best response to x is to make no payments; the outcome is that bill X is passed.
- If $m_x = V_Y$ then both the actions in the previous two cases are best responses by group Y to x.

Buying votes

Subgame perfect equilibria

We conclude that group Y's strategy in a subgame perfect equilibrium has the following properties:

- After a history x for which $m_x < V_Y$, group Y matches group X's payments to the μ legislators to whom X's payments are smallest.
- After a history x for which $m_x > V_Y$, group Y makes no payments.
- After a history x for which $m_x = V_Y$, group Y either makes no payments or matches group X's payments to the μ legislators to whom X's payments are smallest.

Buying votes

Subgame perfect equilibria

Given the properties of group Y's subgame perfect equilibrium strategy, what should X do?

- If it chooses a list of payments x for which $m_x < V_Y$ then group Y matches its payments to a bare majority of legislators, and bill Y passes.
- If it reduces all its payments, the same bill is passed.
- Thus the only list of payments x with $m_x < V_Y$ that may be optimal is $(0, \ldots, 0)$.
- If it chooses a list of payments x with $m_x > V_Y$ then group Y makes no payments, and bill X passes.
- If it reduces all its payments a little (keeping the payments to every bare majority greater than V_Y), the outcome is the same.
- Thus no list of payments x for which $m_x > V_Y$ is optimal.

Buying votes

Subgame perfect equilibria

Conclusion: In any subgame perfect equilibrium we have

1. either $x = (0, \ldots, 0)$ (group X makes no payments)
2. or $m_x = V_Y$ (the smallest sum of group X's payments to a bare majority of legislators is V_Y).

Under what conditions does each case occur?

Buying votes

Subgame perfect equilibria

- If group X needs to spend more than V_X to deter group Y from matching its payments to a bare majority of legislators, then its best strategy is to make no payments ($x = (0, \ldots, 0)$).
- How much does it need to spend to deter group Y? It needs to pay more than V_Y to every bare majority of legislators, so it needs to pay each legislator more than V_Y/μ in which case its total payment is more than kV_Y/μ.
- Thus if $V_X < kV_Y/\mu$, group X is better off making no payments than getting bill X passed by making payments large enough to deter group Y from matching its payments to a bare majority of legislators.

Buying votes

Subgame perfect equilibria

- If $V_X > kV_Y/\mu$, group X can afford to make payments large enough to deter group Y from matching.
- In this case its best strategy is to pay each legislator V_Y/μ, so that its total payment to every bare majority of legislators is V_Y.
- Given this strategy, group Y is indifferent between matching group X's payments to a bare majority of legislators and making no payments.
- The game has no subgame perfect equilibrium in which group Y matches (the argument is similar to the argument that the ultimatum game has no subgame perfect equilibrium in which person 2 rejects the offer 0).
- Thus in any subgame perfect equilibrium group Y makes no payments in response to group X's strategy.

Buying votes

Subgame perfect equilibria

Summing up: if $V_X = kV_Y/\mu$ then the game has a unique subgame perfect equilibrium, in which group Y's strategy is

- match group X's payments to the μ legislators to whom X's payments are smallest after a history x for which $m_x < V_Y$, and
- make no payments after a history x for which $m_x \geq V_Y$,

and group X's strategy depends on the relative sizes of V_X and V_Y:

- if $V_X < kV_Y/\mu$ then group X makes no payments;
- if $V_X > kV_Y/\mu$ then group X pays each legislator V_Y/μ.

If $V_X < kV_Y/\mu$ then the outcome is that neither group makes any payment, and bill Y is passed; if $V_X > kV_Y/\mu$ then the outcome is that group X pays each legislator V_Y/μ, group Y makes no payments, and bill X is passed. (If $V_X = kV_Y/\mu$ then the analysis is more complex.)

Buying votes

Subgame perfect equilibria

Features of the subgame perfect equilibrium:

1. The outcome favors the secon-mover in the game (group Y): group X manages to get the bill X passed only if $V_x > kV_y/\mu$, which is close to $2V_Y$ when k is large.
2. Group Y never makes any payments! According to its equilibrium strategy it is prepared to make payments in response to certain strategies of group X, but given group X's *equilibrium* strategy, it spends not a cent
3. If group X makes any payments (as it does in the equilibrium for $V_X > kV_Y/\mu$) then it makes a payment to every legislator.

Ticktacktoe and Chess

Ticktacktoe, chess, and related games may be modeled as extensive games with perfect information.

- A history is a sequence of moves and each player prefers to win than to tie than to lose.
- Both ticktacktoe and chess may be modeled as finite games, so each game has a subgame perfect equilibrium.
- Both games are strictly competitive games: in every outcome, either one player loses and the other wins, or the players draw. For such games all Nash equilibria yield the same outcome. Further, a player's Nash equilibrium strategy yields at least her equilibrium payoff, regardless of the other players' strategies.
- Because any subgame perfect equilibrium is a Nash equilibrium, the same is true for subgame perfect equilibrium strategies.

Ticktacktoe and Chess

We conclude that in ticktacktoe and chess:

1. either one of the players has a strategy that guarantees she wins, or
2. each player has a strategy that guarantees at worst a draw.

- In ticktacktoe we know that (2) is true.
- Chess is more subtle: it is not known whether White has a strategy that guarantees it wins, or Black has a strategy that guarantees it wins, or each player has a strategy that guarantees at worst a draw.
- The empirical evidence suggests that Black does not have a winning strategy, but this result has not been proved.
- When will a subgame perfect equilibrium of chess be found?

Further reading

- Martin J. Osborne: An Introduction to Game Theory, **Chapters 5, 6 & 7**. Oxford University Press, 2004.